

# DRP Group: Projected Verification Cluster Usage for F17

Jim Bosch, Robert Lupton, John Swinbank, 2017-04-10 (fb3690d)

## Data Release Production

### Datasets

During F17, we expect to continue testing and validation of Data Release Production algorithms primarily by repeated reprocessing of the first HSC Public Data Release (PDR1) on the LSST Verification Cluster (VC).

We expect to perform processing at three different scales:

- The full PDR1 dataset;
- A substantial fraction (nominally 10% of PDR1);
- The HSC “RC” dataset (a subset of PDR1<sup>1</sup> pre-selected for pipeline release testing).

The full PDR1 dataset consists of 6202 exposures, or 17 TB of raw data. It is now available in the /datasets/ filesystem on the VC (see RFC-297, DM-9683). One complete reprocessing of PDR1 requires around 200 TB of storage (see DM-8143); we therefore assume that 10% of PDR1 requires around 20 TB; we expect reprocessing the RC dataset to consume around 7 TB.

Again following DM-8143, we expect one complete reduction of PDR1 to consume around 750 core-weeks of CPU time (and, similarly, 75 core-weeks for a 10% fraction, or 25 core-weeks for the RC dataset). Note that:

- As of April 2017 there are 1152 cores in the VC, so we might reasonably expect that the entire data release can be processed in about 5 days.
- This assumes minimal inefficiency due to workflow; we expect wall-clock time to be rather higher.

## Automated Processing

We expect that some processing takes place automatically, without intervention or explicit request from the DRP team. In each case, processing makes use of the latest weekly release of the LSST stack, with the default configuration; in special circumstances, the DRP team may request an alternative version and/or configuration before the processing run starts.

The pipeline logic will be provided by the DRP team in whatever the currently-accepted standard for LSST data processing is. That is, we expect to use `pipe_drivers/ctrl_pool` style distribution middleware until the point at which a new solution, e.g. one based on SuperTask and Pegasus, becomes available. At that point, the DRP team is responsible for porting their pipelines to the new system.

We expect that regular execution of the relevant pipelines and checking for successful execution will take place outside the scope of DRP. We expect that failures at the execution middleware, hardware or networking layer will be resolved without the need for explicit pipeline intervention. We expect the DRP team to be responsible for triaging and resolving failures in pipeline logic, configuration, etc.

### Note

In the below, we suggest a calendar-based processing scheme. In practice, one which is tied to specific stack releases, rather than to the date, is likely preferable. However, implementing such a scheme would require rethinking the stack release procedure.

## PDR1

To be reprocessed every two months. The results of the last three jobs should be retained: in the steady state this will consume ~600 TB of storage.

## RC Dataset

To be reprocessed weekly. The results of the last four jobs should be retained: in the steady state this will consume ~28 TB of storage.

## Manual Processing

We request a mechanism by which developers may manually trigger processing jobs which will address broadly arbitrary subsets of HSC PDR1 with user specified software versions and configurations, e.g. as supplied through a configuration file (or shell script, etc).

Although DRP developers will be ultimately responsible for the successful execution of these jobs, we request support from NCSA in triaging failures which may be due to cluster or middleware issues.

## Storage

That the total storage requirement for such ad-hoc jobs during F17 will amount to no more than 200 TB. We suggest that this be provisioned in `/project/`, and that it follow the regular policies which apply to that filesystem.

## **Compute**

We expect to consume around 50 core-weeks per calendar week on ad hoc processing (that is, equivalent to two reductions of the RC dataset per week).

## **Calibration Products Production**

### **Datasets**

We expect that data from both TS8 (RFC-301) and the 0.9 m at CTIO (RFC-313) continue to be regularly made available on the /datasets/ filesystem.

On the timescale of F17, we expect these datasets to total no more than 20 TB.

### **Automated Processing**

We do not request any automated processing of data for Calibration Products Producing during F17.

### **Manual Processing**

We expect that developers will manually trigger processing jobs which will address broadly arbitrary subsets of the TS8 & CTIO data with user specified software versions and configurations, e.g. as supplied through a configuration file (or shell script, etc).

Although DRP developers will be ultimately responsible for the successful execution of these jobs, we request support from NCSA in triaging failures which may be due to cluster or middleware issues.

### **Storage**

That the total storage requirement for such ad-hoc jobs during F17 will amount to no more than 50 TB. We suggest that this be provisioned in /project/, and that it follow the regular policies which apply to that filesystem.

## **Compute**

We expect to consume no more than 25 core-weeks per calendar week processing this data.

---

<sup>1</sup>In fact, I believe the existing RC dataset is not, in fact, all public. However, it should be straightforward to define a new RC-sized dataset which is.